

United States Patent Application

Title of the Invention

A NETWORK ROUTING APPARATUS

Inventors

Kazuo SUGAI,

Nobuhito MATSUYAMA.

09767707.012401

- 1 -

A NETWORK ROUTING APPARATUS

BACKGROUND OF THE INVENTION

05676766
The present invention relates to a network routing apparatus. It particularly relates to a network routing apparatus which is adapted to a network routing apparatus such as a router, a repeater, or the like, in a computer network system and which can retrieve a forwarding address of an input packet at a high speed.

A network routing apparatus such as a router, a bridge, or the like, is generally used for connecting a plurality of networks to one another in a network system. The router examines a destination address of a packet received from a connected network to thereby determine the forwarding address of the packet and forwards the packet to a network to which a router or host of the forwarding address is connected.

Whenever the router receives a packet from a network interface board, the router determines a packet output interface on the basis of the destination information contained in header information of the packet, further determines a forwarding host or router out of hosts or routers connected to the packet output interface, and forwards the packet to the determined host or router. A higher-functional router performs, in addition to the function of forwarding the packet, a function (filtering function) for aborting a packet in

accordance with the header conditions of the packets, a function (priority control function) for determining output priority of a plurality of packets to output the packets in the order of priority of the plurality of
5 packets when such packets are to be forwarded, a function for rewriting Type of Service (ToS) contained in the header information of the packet, and so on.

Generally, the router further has a statistical function for counting the number of packets
10 forwarded, the number of bytes forwarded, the number of packets aborted by filtering in the router or by traffic congestion of packets in the router, the number of (non-routed) packets undefined in terms of output destination as a result of retrieval (routing
15 retrieval) of output addresses of packets, the number of packets addressed to the router itself, the number of packets satisfying the aborting condition of filtering, and so on. These pieces of statistical information are displayed on a terminal for managing
20 router operations or router status indication, or the like, or these pieces of statistical information are sent to a network management apparatus in a network so as to be used for management of the network.

Functions other than the statistical function
25 may be required as follows, in the case where the number of packet data transfer bytes needs to be measured. When, for example, a network provider receives packets from users, the packets are classified

09767707-012401

into several groups (hereinafter referred to as "flows") (for example, according to source/destination sub-network, TCP/IP port number, address, etc.) on the basis of the header conditions of the packets. The data transfer rate of each flow is monitored whenever this classification is made. When the data transfer rate in this flow is too high, packets sent from the users in the flow over the bandwidth on the basis of contracts between the network provider and the users may be aborted or preferentially aborted at the time of traffic congestion. This function is called "bandwidth policing function". Further, when network users transmit packets to a network of a network provider, the output rate of packet data belonging to the flow of the network provider may be limited, by a router, within the bandwidth on the basis of a contract between the router and the network provider. This function is called "shaping function". To achieve the bandwidth policing function or the shaping function, the number of data transfer bytes (data transfer rate) of packets flowing in the router needs to be measured at intervals of a predetermined time.

Processing in a router is classified into the aforementioned process which needs to be performed whenever a packet enters the router, a process in which routing information in a network system is transferred between routers so that a routing table is generated in the router, and a process which need not be performed

09767707 012401
104210 202260

whenever a packet such as statistical information of the router transferred between the router and a network management apparatus enters the router.

09767707-012401

The aforementioned processing in a router is heretofore carried out by executing software processing. To improve performance of processing required whenever a packet enters a router, a plurality of units (RPs: Routing Processors) for performing this processing are provided in the router. Network interface boards are connected to the RPs correspondingly and respectively. The RPs are connected to one another through a bus or through a cross bar switch so that packets are forwarded between network interface boards connecting individual RPs, respectively. In this manner, high-speed processing in the router is attained. An example of the aforementioned router is an apparatus that is disclosed in U.S. Patent No. 5,434,863 (JP-A-5-199230).

Traffic flowing in a network has, however, increased so that the routing process needs to be performed at a higher speed. As a result, software processing used for achieving the routing process has been replaced by hardware processing.

SUMMARY OF THE INVENTION

25 Traffic flowing in a network has been increasing at a pace higher than improvement of processing capacity achieved with the advance of

hardware element techniques such as LSI manufacturing techniques in recent years. If the router or bridge processes the traffic, through hardware, whenever the router or bridge receives a packet from a network interface board, the router or bridge cannot cope with the increase in network traffic.

09767707.012401
10 Upon such circumstances, an object of the present invention is to provide a routing apparatus in which hardware units for performing a routing process are arranged in parallel to one another to thereby improve the total performance of the router more greatly than the processing capacity achieved with the advance of hardware element techniques such as LSI manufacturing techniques so that the routing apparatus
15 can sufficiently cope with the increase in network traffic occurring in recent years.

According to the present invention, the router has network interface boards (NIFs) for performing a packet I/O interface operation, routing
20 processors (RPs) for retrieving forwarding addresses of packets on the basis of header information of the packets and for forwarding the packet to the forwarding addresses, a routing manager (RM) for performing control of the respective RPs and transmission/
25 reception of control packets such as routing protocols, and a cross bar switch (CSW) for connecting the respective RPs to the RM and for performing forwarding of packets between RPs and between RM and RPs. A

plurality of packet forwarding units for retrieving destination addresses of packets on the basis of header information of the packets and for forwarding the packets to the destination addresses are provided in
5 each of the RPs. Packets sent from the NIFs and packets sent from the CSW are distributed one by one to the packet forwarding units.

To limit the process to be performed by the packet forwarding units to thereby simplify the
10 configuration of the packet forwarding units, the packet forwarding units may be configured so that packet forwarding units for forwarding packets from the NIF to the CSW are provided separately from packet forwarding units for forwarding packets from the CSW to
15 the NIF. To shorten the idle time of the packet forwarding units to thereby improve the processing performance, per packet forwarding unit, of the router, the packet forwarding units may be configured so that each of the packet forwarding unit serves as a unit for
20 forwarding packets from NIF to CSW and for forwarding packets from CSW to NIF.

Further, according to the present invention, each of the RPs has forward packet distribution units so that packets sent from the NIF or from the CSW are
25 distributed to the plurality of packet forwarding units in the RP. The forward packet distribution units distribute packets by a round-robin method, by an out-of-order method, or by a hash method. The round-robin

09767707.012401

method is a method in which packets are distributed sequentially to the packet forwarding units in entering order of the packets. The out-of-order method is a method in which packets are distributed to empty ones
5 of the packet forwarding units. The hash method is a method in which packet forwarding units used for distributing packets are univocally determined on the basis of the header conditions of the packets.

A function (hash function) used in the hash
10 method for determining forwarding addresses of packets can be set in such a condition that a series of packets to be continuously sent from a packet sender address to a forwarding address are all assigned to one packet forwarding unit. A pseudo-random number sequence
15 generated on the basis of the header conditions may be used as the hash function so that packets can be assigned to the packet forwarding units as uniformly as possible.

Further, according to the present invention,
20 each of the RPs has forward packet rearrangement units for rearranging the packets subjected to the forwarding process by the plurality of packet forwarding units in the RP and for sending out the rearranged packets to the NIF or to the CSW.

25 When a round-robin method or an out-of-order method is used as a distributing method in the forward packet distribution units, the forward packet rearrangement units rearrange packets in the order of

09757707.012401
T04270.202960

the packets entering the router after each of the packet forwarding units performs a packet forwarding process because the order of completion of the packet forwarding process in the packet forwarding units in the RP may be reversed to the entering order of the series of packets to be continuously sent from the packet sender address to the destination address. When an out-of-order method is used as a distributing method in the forward packet distribution units, packets are numbered sequentially when the forward packet distribution units assign the packets to the packet forwarding units so that the forward packet rearrangement units can rearrange the packets in the order of the packets entering the router. In this case, the packets are taken around with the sequence numbers in the packet forwarding units, so that the forward packet rearrangement units rearrange the packets in the order of the sequence numbers assigned to the packets, and send out the rearranged packets.

Further, according to the present invention, each of the packet forwarding units has a forward packet processing queue waiting for the start of the packet forwarding process after reception of packets from the forward packet distribution units, and a forward packet rearrangement queue waiting for sending of the packets to the forward packet rearrangement units after the completion of the packet forwarding process in the packet forwarding units.

09767707.012401

09767707, 012401

A plurality of packet retrieving units for retrieving packet header information are provided in each packet forwarding unit so that packet header retrieving processes among the processes performed by the packet forwarding unit are arranged in parallel to one another to attain high-speed processing. A packet header operating mechanism for extracting the headers of packets is provided in the packet forwarding unit so that header information of the packets can be extracted. A retrieval packet distribution unit for distributing the extracted packet headers to the respective packet retrieving units is provided in the packet forwarding unit. Each of the packet retrieving units retrieves forwarding addresses of packets on the basis of the header information of the packets, retrieves the packets to be aborted or to be forwarded (filtering retrieval), and retrieves output priority of packets (communication quality retrieval). Each of the packet retrieving units returns information concerning the forwarding addresses of the packets, information as to whether the packets are to be aborted or not and information concerning the output priority of packets. A retrieval packet rearrangement unit for rearranging retrieval results in the packet retrieving units and for transferring the rearranged retrieval results to the packet header operating mechanism is further provided in the packet forwarding unit.

To attain higher-speed processing, parallel

arrangement of the packet forwarding units and parallel arrangement of the packet retrieving units may be performed simultaneously.

Further, according to the present invention,
5 statistical information picked up in the plurality of packet forwarding units or in the plurality of packet retrieving units is tabulated as follows. A software program managing the router as a whole reads statistical information in the respective packet
10 forwarding units or in the respective packet retrieving units and tabulates the information by software processing so that the tabulated information is displayed as total statistical information of the router on a management terminal or so that the
15 tabulated information is sent to the network management apparatus. Alternatively, a processor is provided in each of the RPs. A software program operating on the processor reads statistical information in the respective packet forwarding units or in the respective
20 packet retrieving units and tabulates the statistic information by software processing.

Further, according to the present invention, packets are classified into several groups (hereinafter referred to as "flows") in accordance with the header
25 conditions of the packets. Data transfer rate of each flow is measured whenever this classification is made. The flows are numbered so that the packets can be aborted in accordance with the data transfer rate or

09767707-012401

aborted easily at the time of traffic congestion or so that the packet sending rate can be limited at the time of sending of the packets. The flow numbers are outputted from the respective packet forwarding units
5 or from the respective packet retrieving units. The data transfer rate in every flow is measured in the forward packet rearrangement units or in the retrieval packet rearrangement unit in accordance with the flow numbers.

10 BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a diagram showing the overall configuration of a network system using routers according to an embodiment of the present invention;

Fig. 2 is a block diagram showing the overall
15 internal configuration of a router according to an embodiment of the present invention;

Fig. 3 is a block diagram showing the internal configuration of a routing processor having parallel packet forwarding units according to an
20 embodiment of the present invention (in the case where the packet forwarding units are independent of one another on the basis of directions of transmission thereof);

Fig. 4 is a block diagram showing the
25 internal configuration of a routing processor having parallel packet forwarding units according to another embodiment of the present invention (in the case where

09767707.012401
TOH270 20229260

one packet forwarding unit performs a bilateral packet forwarding process);

Fig. 5 is a block diagram showing the internal configuration of a packet forwarding unit according to an embodiment of the present invention;

Fig. 6 is a block diagram showing the internal configuration of a packet forwarding unit and a packet forwarding mechanism contained in the packet forwarding unit in the case where packet retrieving units are not arranged in parallel to one another;

Fig. 7 is a block diagram showing the internal configuration of a packet forwarding unit and a packet forwarding mechanism contained in the packet forwarding unit in the case where packet retrieving units are arranged in parallel to one another;

Fig. 8 is a block diagram showing the internal configuration of a packet retrieving unit in an embodiment of the present invention;

Fig. 9 is a conceptual diagram showing a process in which statistical information picked up in packet retrieving units is read and tabulated in a routing processor;

Fig. 10 is a block diagram showing the internal configuration of a packet forwarding mechanism which shows an extracted portion in which a retrieval packet rearrangement unit measures a transfer rate of the flow in accordance with each of flow numbers retrieved by a plurality of packet retrieving units;

09757707.012401

Fig. 11 is a time chart showing packet forwarding processing of packet forwarding units in the case where a round-robin method is used as a distributing method in a forward packet distribution

5 units;

Fig. 12 is a time chart showing packet forwarding processing of packet forwarding units in the case where an out-of-order method is used as a distributing method in a forward packet distribution

10 units;

Fig. 13 is a time chart showing packet forwarding processing of packet forwarding units in the case where a hash method is used as a distributing method in a forward packet distribution units;

15 Fig. 14 is a diagram showing a flow retrieval table;

Fig. 15 is a diagram showing a first input circuit limiting system; and

20 Fig. 16 is a diagram showing a second input circuit limiting system.

DESCRIPTION OF THE EMBODIMENTS

Fig. 1 is a diagram showing the overall concept of a network system using routing apparatuses according to the present invention. In Fig. 1, the network system comprises: hosts 1, 2 and 3 which may be packet senders/receivers; routers 4 and 5 for forwarding packets in accordance with destination

09767707-012401
T042T0720729260

addresses written in headers of the respective packets;
and sub-networks 6, 7 and 8 to which the hosts and the
routers are directly connected and in which data are
transferred between the hosts and the routers. The
5 routers 4 and 5 may be also packet senders/receivers.

When, for example, a packet is to be
forwarded from the host 1 to the host 2, the host 1
forwards the packet from an interface 10 to the router
4 connected to the interface 10 through the sub-network
10 6. The router 4 receives the packet from an interface
40 connected to the sub-network 6 and forwards the
packet from an interface 41 to the router 5 connected
to the interface 41 through the sub-network 7.
Similarly, the router 5 receives the packet from an
15 interface 50 and forwards the packet from an interface
51 to the host 2 connected to the interface 51 through
the sub-network 8. The host 2 receives the packet from
an interface 20. Addresses in sub-networks are used
for specifying routers or hosts peculiar to the sub-
20 networks.

Fig. 2 is a block diagram showing the overall
internal configuration of a router. In Fig. 2, the
router 4 includes: network interface boards (NIFs) 450
to 455 for performing packet I/O interface operations;
25 routing processors (RPs) 440 to 442 for retrieving a
forwarding address of a packet on the basis of header
information of the packet and forwarding the packet to
the forwarding address; a routing manager (RM) 42 for

09767707 012401
T04210 2029260

performing general management of the router, control of the respective RPs, management of routing information, etc. on the basis of transmission/reception of a control packet such as a routing protocol; and a cross
5 bar switch (CSW) 43 for connecting the respective RPs and the RM to perform packet transfer between the RPs and between the RM and the RPs. The NIFs 450 to 455 are connected to network interfaces 400 to 405 respectively. The network interfaces 400 to 405 in
10 Fig. 2 correspond to the network interfaces 40 and 41 in Fig. 1. Such network interfaces are classified into LAN (Local Area Network) such as Ethernet, or the like; WAN (Wide Area Network); ATM (Asynchronous Transfer Mode); etc. The NIFs 450 to 455 are constituted by
15 different hardware in accordance with the interface classification.

Figs. 3 and 4 are block diagrams showing the internal configuration of the RP 440. Figs. 3 and 4 show two examples of the internal configuration of the
20 routing processor according to the present invention in the case where packet forwarding units are arranged in parallel to one another. The packet forwarding units 4400 to 4407 (Fig. 3) and 4410 to 4413 (Fig. 4) perform a process for retrieving a destination address of a
25 packet on the basis of header information of the packet in the RP 440 and for forwarding the packet to the destination address. To distribute load, a plurality of packet forwarding units are provided in the RP 440

so that packets sent from the NIF and CSW are assigned one by one to the plurality of packet forwarding units. Fig. 3 shows an example of configuration of the RP 440 in the case where packet forwarding units 4400 to 4403
5 for forwarding packets sent from the NIF are provided separately from packet forwarding units 4404 to 4407 for forwarding packets sent from the CSW. Fig. 4 shows an example of configuration of the RP 440 in the case where the packet forwarding units 4410 to 4413 serve as
10 units for forwarding packets sent from the NIF and for forwarding packets sent from the CSW.

Forward packet distribution units 4420 and 4422 (Fig. 3) and 4430 and 4432 (Fig. 4) perform a process by which a packet sent from the NIF or CSW is
15 distributed to any one of the plurality of packet forwarding units. A round-robin method, an out-of-order method, a hash method, or the like, is used as a packet distributing method. The round-robin method is a method in which packets are distributed to the packet
20 forwarding units in the entering order. The out-of-order method is a method in which packets are distributed to empty ones of the packet forwarding units. The hash method is a method in which a packet forwarding unit to be assigned for a packet is
25 determined univocally on the basis of the header condition of the packet. Fig. 11 is a time chart showing a packet forwarding process in each of the packet forwarding units in the case where a round-robin

09757707.012401
T042707.2022960

method is used as the packet distributing method in the forward packet distribution units. Fig. 12 is a time chart showing a packet forwarding process in each of the packet forwarding units in the case where an out-
5 of-order method is used as the packet distributing method in the forward packet distribution units. Fig. 13 is a time chart showing a packet forwarding process in each of the packet forwarding units in the case where a hash method is used as the packet distributing
10 method in the forward packet distribution units. In each of the time charts of Figs. 11, 12 and 13, the numeral references show a sequence of packets input into the RP 440, and the length of each rectangle shows the time required for the packet forwarding unit to
15 process the packet. Each of Figs. 11, 12 and 13 shows the case where packets long in processing time enter as odd-numbered packets and packets short in processing time enter as even-numbered packets.

For example, a function (hash function) used
20 in the hash method to determine a forwarding address of a packet is provided with a condition that a series of packets to be continuously sent from a sender address to a forwarding address must be all distributed to one packet forwarding unit. Under this condition, the
25 series of packets to be continuously sent from the sender address to the forwarding address can be delivered from the sender address to the forwarding address in the entering order after the packets are

09767707-012401
T042707-2072960

sent in the order of finishing of the packet forwarding process even in the case where the time required for each packet forwarding process varies.

Moreover, a pseudo-random number sequence
5 generated on the basis of the packet header condition may be used as the hash function so that packets can be distributed to the packet forwarding units as uniformly as possible. As an example of the hash function in the case where two packet forwarding units are provided,
10 there is a method of exclusively ORing all bits of a 32-bit source IP address and all bits of a 32-bit destination IP address. In this case, to which one of the two forwarding units the packet is to be assigned is determined on the basis of one bit obtained by
15 exclusive ORing of all bits. When four packet forwarding units are provided, a hash function capable of outputting not one bit but two bits is used. Fig. 13 shows the case where the first and third packets are assigned to the packet forwarding unit 4410, the
20 second and fourth packets are assigned to the packet forwarding unit 4411, the fifth and seventh packets are assigned to the packet forwarding unit 4412, and the sixth and eighth packets are assigned to the packet forwarding unit 4413.

25 When a round-robin method or an out-of-order method is used as a distributing method in the forward packet distribution units 4420 and 4422 (Fig. 3) and 4430 and 4432 (Fig. 4), the forward packet

09767707 012401
TOP SECRET

rearrangement units 4421 and 4423 (Fig. 3) and 4431 and 4433 (Fig. 4) rearrange packets in the order of the packets entering the router 4 after the packet forwarding units 4400 to 4407 (Fig. 3) and 4410 to 4413 (Fig. 4) have finished the packet forwarding processes. This order rearrangement is made because the order of completion of the packet forwarding process in the packet forwarding units 4400 to 4407 (Fig. 3) and 4410 to 4413 (Fig. 4) in the RP may be reversed to the entering order of the series of packets to be continuously sent from the sender address to the destination address. When an out-of-order method is used as a distributing method in the forward packet distribution units 4420 and 4422 (Fig. 3) and 4430 and 4432 (Fig. 4), packets are numbered sequentially when the forward packet distribution units 4420 and 4422 (Fig. 3) and 4430 and 4432 (Fig. 4) assign the packets to the packet forwarding units 4400 to 4407 (Fig. 3) and 4410 to 4413 (Fig. 4) so that the forward packet rearrangement units 4421 and 4423 (Fig. 3) and 4431 and 4433 (Fig. 4) can rearrange the packets in the order of the packets entering the router 4. In this case, the packets are taken around with the sequence numbers in the packet forwarding units 4400 to 4407 (Fig. 3) and 4410 to 4413 (Fig. 4), so that the forward packet rearrangement units 4421 and 4423 (Fig. 3) and 4431 and 4433 (Fig. 4) rearrange the packets in the order of the sequence numbers assigned to the packets, and send out

09767707, 012401

the rearranged packets.

When a hash method is used as a distributing method in the forward packet distribution units 4420 and 4422 (Fig. 3) and 4430 and 4432 (Fig. 4), the packet output order may be changed or disordered among the packets sent to any one of the packet forwarding units 4400 to 4407 (Fig. 3) and 4410 to 4413 (Fig. 4). In this case, the corresponding one of the forward packet rearrangement units 4421 and 4423 (Fig. 3) and 4431 and 4433 (Fig. 4) may send out the packets in the order of termination of the packet forwarding process. When the packet forwarding process is terminated in the plurality of packet forwarding units 4400 to 4407 (Fig. 3) and 4410 to 4413 (Fig. 4), the packet forwarding units 4400 to 4407 (Fig. 3) and 4410 to 4413 (Fig. 4) send out the packets in the order of termination of the packet forwarding process.

Next, Fig. 5 is a block diagram showing the internal configuration of a packet forwarding unit in the present invention.

The packet forwarding units 4400 to 4407 (Fig. 3) or 4410 to 4413 (Fig. 4) have similar ways to have queues waiting for internal processing though they are different in the packet forwarding direction, that is, they are directed in the direction from the NIF 450 to the CSW 43 or from the CSW 43 to the NIF 450. Fig. 5 is a block diagram showing the way that the packet forwarding unit 4400 as a representative has internal

09767707 012401
F012707 2025/6/6

processing queues waiting for the packet processing.

In Fig. 5, the packet forwarding unit 4400 has a packet forwarding mechanism 44000 for executing a packet forwarding process, a forward packet processing queue 5 44001 waiting for the start of packet processing in the packet forwarding mechanism 44000 after assignment of packets to the packet forwarding unit 4400 by the forward packet distribution unit 4420, or the like, and a forward packet rearrangement queue 44002 waiting for 10 the rearrangement and output of the packets by the forward packet rearrangement unit 4421, or the like, after the completion of processing in the packet forwarding mechanism 44000.

The forward packet processing queue 44001 and 15 the forward packet rearrangement queue 44002 absorb the deviation of the packet processing time in the packet forwarding unit 4400, or the like, and absorb the deviation of the time required for distributing packets to the packet forwarding unit 4400, or the like, if a 20 hash method is used as a distributing method in the forward packet distribution unit 4420, or the like. There is an effect that the idle time of the packet forwarding unit 4400, or the like, is shortened.

An embodiment will be described below in the 25 case where packet retrieving portions in the packet forwarding unit 4400 are arranged in parallel to one another in order to attain high-speed processing. Fig. 6 is a block diagram of the internal configuration of

09767707-012401
1042707-2029260

the packet forwarding unit 4400 and the packet forwarding mechanism 44000 contained in the packet forwarding unit 4400 in the case where packet retrieving units for retrieving the packet headers are not arranged in parallel. Fig. 7 shows the diagram of the internal configuration of the packet forwarding unit 4400 and the packet forwarding mechanism 44000 contained in the packet forwarding unit 4400 in the case where packet retrieving units are arranged in parallel to one another.

In Fig. 6, the packet forwarding mechanism 44000 has a packet header operating mechanism 440000, and a packet retrieving unit 440001. The packet header operating mechanism 440000 extracts header portions of packets and delivers the header portions to the packet retrieving unit 440001. The packet retrieving unit 440001 retrieves forwarding addresses of the packets on the basis of the packet header information sent from the packet header operating mechanism 440000, retrieves the packets to be aborted or to be forwarded (filtering retrieval), retrieves output priority of the packets (communication quality retrieval), and so on. The packet retrieving unit 440001 delivers information concerning the forwarding addresses of the packets, information as to whether the packets are to be aborted or not, information concerning the output priority of the packets, and so on, to the packet header operating mechanism 440000. The packet header operating

09767707.012401
TOTAL 202926

mechanism 440000 forwards or aborts the packets on the basis of the information delivered by the packet retrieving unit 440001, such as information concerning the forwarding addresses of the packets and information
5 as to whether the packets are to be aborted or to be forwarded.

Fig. 7 shows the configuration of the packet forwarding unit 4400 in the case where packet retrieving units are arranged in parallel to one
10 another. The packet forwarding unit 4400 has a packet forwarding mechanism 44000. The packet forwarding mechanism 44000 has a packet header operating mechanism 440000, packet retrieving units 440001 to 440004 each equivalent to the packet retrieving unit 440001 shown
15 in Fig. 6, a retrieval packet distribution unit 440005 for assigning a packet header retrieving process to one of the packet retrieving units 440001 to 440004 in the packet forwarding mechanism 44000, and a retrieval packet rearrangement unit 440006 for rearranging
20 retrieval results of the packet retrieving units 440001 to 440004.

A round-robin method, an out-of-order method or a hash method can be used as a packet retrieving unit distributing method in the retrieval packet
25 distributing unit 440005 in the same manner as that in the forward packet distribution units 4420 and 4422 (Fig. 3) and 4430 and 4432 (Fig. 4). The characteristic of each distributing technique is the

0076707 012401

same as that used in the forward packet distribution units. The same rearranging method as that used in the forward packet rearrangement units 4421 and 4423 (Fig. 3) and 4431 and 4433 (Fig. 4) is applied to rearrangement in the retrieval packet rearrangement unit 440006.

Fig. 8 shows the diagram of the internal configuration of a packet retrieving unit in the present invention.

10 The packet retrieving units 440001 to 440004 shown in Fig. 7 are equal in the way that each of the units has queues waiting for internal processing. Fig. 8 is a block diagram showing the way that the packet retrieving unit 440001 as a representative has queues waiting for internal processing. In Fig. 8, the packet retrieving unit 440001 has a packet retrieving mechanism 4400010 for executing a packet retrieving process, a retrieval packet queue 4400011 waiting for the start of packet processing in the packet retrieving mechanism 4400010 after assignment of packets to the packet retrieving unit 440001, or the like, by the retrieval packet distribution unit 440005, and a retrieval packet rearrangement queue 4400012 waiting for the rearrangement and output of packets in the retrieval packet rearrangement unit after completion of processing in the packet retrieving mechanism.

20
25

The retrieval packet processing queue 4400011 and the retrieval packet rearrangement queue 4400012

09767707-012401
T042T0-20229260

absorb the deviation of the packet processing time in the packet retrieving unit 440001, or the like, and absorb the deviation of the time required for distributing packets to respective packet retrieving
5 units when a hash method is used as a distributing method in the retrieval packet distribution unit. There is an effect that the idle time of the respective packet retrieving units is shortened.

A method of picking up statistical
10 information will be described below in the case where the aforementioned packet forwarding units 4400 to 4407 (Fig. 3) or 4410 to 4413 (Fig. 4) or the packet retrieving units 440001 to 440004 are arranged in parallel to one another. For example, in the packet
15 forwarding units 4400 to 4407 (Fig. 3) or in the packet forwarding units 4410 to 4413 (Fig. 4), the number of packets and the number of bytes in packets subjected to the forwarding process and the number of packets aborted by filtering in the router or by traffic
20 congestion of packets in the router are counted. In the packet retrieving units 440001 to 440004, the number of (non-routed) packets undefined in terms of output address, the number of packets addressed to the router itself, the number of packets satisfying the
25 aborting condition owing to filtering, and so on, are counted as a result of retrieval (routing retrieval) of packet output address.

In the present invention, the packet

09767707.012401
204420/20229266

forwarding units 4400 to 4407 (Fig. 3), the packet forwarding units 4410 to 4413 (Fig. 4) or the packet retrieving units 440001 to 440004 are arranged in parallel to one another. Hence, to pick up the

5 aforementioned statistical information in the packet forwarding units 4400 to 4407 (Fig. 3), the packet forwarding units 4410 to 4413 (Fig. 4) or the packet retrieving units 440001 to 440004, a software program operating on the RM 42 shown in Fig. 2, or a processor

10 is provided in each of the RPs 440 to 442 shown in Fig. 2 so that a software program is operated on the processor. The software program read statistical information picked up in the packet forwarding units 4400 to 4407 (Fig. 3), the packet forwarding units 4410

15 to 4413 (Fig. 4) or the packet retrieving units 440001 to 440004 so that the statistical information picked up in all the parallel packet forwarding units 4400 to 4407 (Fig. 3), all the parallel packet forwarding units 4410 to 4413 (Fig. 4) or all the parallel packet

20 retrieving units 440001 to 440004 is added up.

Fig. 9 is a conceptual diagram showing a process in which statistical information picked up in the packet retrieving units in the present invention is read and tabulated in the routing processor.

25 In Fig. 9, the packet retrieving units 440001 to 440004 in the packet forwarding mechanism 44000 have statistical counters 4400013 to 4400043 respectively. The values of the statistical counters 4400013,

4400023, 4400033 and 4400043 picked up in the packet
retrieving units 440001 to 440004 respectively are read
and tabulated by the processor 445 in the RP 440
through a register access bus. Portions not directly
5 related to collection of packets are not shown in Fig.
9.

Further, functions other than the statistical
function may be required as follows in the case where
the number of packet data transfer bytes needs to be
10 measured. When, for example, a network provider
receives packets from users, the packets are classified
into several groups (hereinafter referred to as
"flows") on the basis of the header conditions of the
packets. Data transfer rate of each flow is monitored
15 whenever this classification is made. When the data
transfer rate in this flow is too high, packets sent
from the users in the flow over the bandwidth on the
basis of contracts between the network provider and the
users may be aborted or preferentially aborted at the
20 time of traffic congestion. This function is called
"bandwidth policing function". Further, when network
users transmit packets to a network of a network
provider, the output rate of packet data belonging to
the flow of the network provider may be limited, by a
25 router, within the bandwidth on the basis of a contract
between the router and the network provider. This
function is called "shaping function". To achieve the
bandwidth policing function or the shaping function,

09767707, 012401
10426 202926

the number of data transfer bytes (data transfer rate) of packets flowing in the router needs to be measured at intervals of a predetermined time.

When a plurality of packet forwarding units
5 4400 to 4407 (Fig. 3), a plurality of packet forwarding units 4410 to 4413 (Fig. 4) or a plurality of packet retrieving units 440001 to 440004 are arranged in parallel to one another, one flow is distributed and processed to the plurality of packet forwarding units
10 4400 to 4407 (Fig. 3), the plurality of packet forwarding units 4410 to 4413 (Fig. 4) or the plurality of packet retrieving units 440001 to 440004. Therefore, to obtain the total data transfer rate of one flow, the data transfer rate of the flow in the
15 respective packet forwarding units 4400 to 4407 (Fig. 3), the respective packet forwarding units 4410 to 4413 (Fig. 4) or the respective packet retrieving units 440001 to 440004 needs to be picked up at a common point. To pick up the data transfer rate of one flow
20 at a common point, the flows are numbered. As one of packet header retrieval results, the flow numbers are outputted from the respective packet forwarding units 4400 to 4407 (Fig. 3), the respective packet forwarding units 4410 to 4413 (Fig. 4) or the respective packet
25 retrieving units 440001 to 440004 so that the data transfer rate of the flow in accordance with the flow numbers is measured in the forward packet rearrangement unit 4421, 4423, 4431 or 4433 or the retrieval packet

09767707 012401
T04210 20229260

rearrangement unit 440006. When the bandwidth policing function is achieved, packets are aborted or packets are aborted easily in accordance with the data flow rate of the flow (if, for example, the data transfer rate of the flow reaches or exceeds a threshold). Furthermore, traffic statistical information can be obtained in accordance with the flow numbers. Further, when the shaping function is achieved, the output is limited by the router so that the data transfer rate of the flow does not exceed the bandwidth on the basis of a contract between the network provider and the router.

Fig. 10 is a block diagram showing the internal configuration of a packet forwarding mechanism which shows an extracted portion in which a retrieval packet rearrangement unit measures the respective data transfer rate of the flow in accordance with each of flow numbers retrieved by a plurality of packet retrieving units in the present invention. The retrieval packet rearrangement unit 440006 has a transfer rate counter 4400060 which is provided for each flow numbers so that information concerning the packet data transfer rate for each flow such as the number of packets and the number of bytes for each flow is stored in the transfer rate counter 4400060 in accordance with the flow number. The packet retrieving units 440001 to 440004 output information containing the flow number as one of packet retrieval results. The retrieval packet rearrangement unit 440006 makes

the transfer rate counter 4400060 measure the packet data transfer rate of each flow in accordance with the flow numbers. Portions not directly related to the measurement of the packet data transfer rate of each flow are not shown in Fig. 10.

Flow retrieval will be described below. Fig. 14 is a view for explaining a flow retrieval table.

As shown in Fig. 14, for example, the flow retrieval table contains, as a comparison field 101, sender IP address, destination IP address, packet length, IP priority, IP high-order protocol, delivery confirmation flag, forwarding TCP/UDP port, and designation TCP/UDP port, and contains, as an action field 102, filtering (pass/abortion), tunneling (encapsulation/non-encapsulation) and QoS (delay class, abortion class, bandwidth, etc.).

A specific method of QoS flow retrieval will be described below. Paying attention to the QoS flow retrieval, description will be made by way of example. The same rule can be applied to various kinds of flow retrieval. Incidentally, the aforementioned flow control information may be mixedly stored in the action field 102 or such a flow retrieval table may be provided for each flow.

First, a linear search system will be described below. When QoS control information of a packet is judged as one of actions, entries set in advance are read from an entry table in the descending

09767707.012401

order. Then, the judgement is made as to whether the QoS control information coincides with the value of the header portion of the packet and all the effective flow conditions in the comparison field 101 in this entry or
5 not. When coincidence is obtained, the QoS control information in the action field 102 in this entry is judged as QoS control information of the packet and the QoS flow retrieval is terminated. When coincidence with the flow condition is retrieved, the QoS control
10 information in the action field 102 is decided as QoS control information and the flow retrieval is terminated without executing retrieval in the next entry.

In the aforementioned linear search system,
15 it may be difficult to execute QoS control or filtering at a high speed in a network having a great deal of entries set therein. Therefore, an input circuit limiting system, or the like, may be preferably used because the input circuit limiting system can perform
20 flow retrieval at a high speed in comparison with the linear search system even in the case where a great deal of entries are set. The input circuit limiting system will be described below in brief. In the input circuit limiting system, only the entry coincident with
25 the input circuit number constituting the comparison field of the linear search system is retrieved to attain high-speed processing.

Fig. 15 is a diagram for explaining the first

00767707.012401

input circuit limiting system. In the first input circuit limiting system, an entry 511-i obtained by removing the input circuit number and the input circuit number effective bit from the comparison field of the

5 linear search system is set in accordance with the corresponding input circuit. For example, a flow condition portion 521-i contains an SIP upper limit 501, an SIP lower limit 502, a DIP upper limit 503 and a DIP lower limit 504 as conditions for identifying a

10 sender and a destination user. The flow condition portion 521-I further contains an IP effective bit 562 indicating that the upper and lower bits of SIP and DIP are effective, an SPORT 505 as a sender port, a DPORT 506 as a destination port, a port effective bit 563

15 indicating that the SPORT 505 and the DPORT 506 are effective, and so on. For example, a QoS control information portion 530-i contains QoS control information 507 used with a preferential forwarding function. Because only the entry 511-i coincident with

20 the input circuit number as a flow condition is retrieved, the input circuit number need not be set in the entry 511-i. At the time of flow retrieval, only the entry 511-i assigned to the input circuit to which a packet is input is retrieved.

25 When an entry 511-i not related to an input circuit number is to be set in the first input circuit limiting system (for example, when traffic of Telnet inputted through all input circuits is set as "high

09757707-012401

priority"), it is necessary to set the entry 511-i in the input circuit number (= N). In this case, the efficiency of memory for achieving the entry table may be worsened. Therefore, a higher-speed input circuit limiting system will be described below.

Fig. 16 is a diagram for explaining the second input circuit limiting system. In the second circuit limiting system, a list 540-ij which is an address of an entry table 750 is set in a list table for every input circuit. For example, a list 540-11 having the list table address "1" shows an address of an entry 511-1, and a list 540-12 having the list table address "2" shows an address of an entry 511-H. At the time of flow retrieval, only the lists 540-ij assigned to the corresponding input circuit through which a packet is supplied are read, so that the entry 511-i pointed by one list 540-ij is read. If a list 540-ij having a small bit width (for example, having about 10 bits per 1024 entries) is retained for every input circuit and an entry 511-i having a large bit width is shared to the respective input circuits, memory for achieving the entry table can be used effectively. Hence, a large number of entries 511-i can be set while high-speed processing is achieved.

25 An output circuit limiting system is another embodiment of the flow detecting system. In the output circuit limiting system, only the entry 511-i coincident with the output circuit number as a flow

09767707 012401
TOHOTO 707/9/60

condition is processed in the same manner as that in the aforementioned input circuit limiting system to thereby achieve high-speed flow detection.

There is further an SAMAC limiting system
5 using source addresses SAMAC of an MAC (Media Access Control) layer instead of the input circuit number of the header information as a flow condition. In the SAMAC limiting system, when an SAMAC group is defined as a group of source addresses SAMAC and the entry is
10 limited by an SAMAC identifier which is an identifier of the SAMAC group, flow retrieval can be executed in the same manner as that in the aforementioned input limiting system.

Incidentally, parallel arrangement of packet
15 forwarding units as shown in Figs. 3 and 4 and parallel arrangement of packet retrieving units as shown in Fig. 7 can be executed simultaneously. When the two parallel arrangements are executed simultaneously, the degree of parallelism of packet retrieving units in
20 each RP is obtained by multiplying the degree of parallelism of packet forwarding units and the degree of parallelism of packet retrieving units.

According to the aforementioned embodiment, a plurality of packet forwarding units for retrieving a
25 designation address of packets on the basis of header information of the packets and for forwarding the packets are provided in each RP, and packets sent from the NIFs and packets sent from the CSW are distributed

0975707 012401

one by one to the packet forwarding units. Hence, the packet forwarding process can be performed at a high speed in accordance with the degree of parallelism of the packet forwarding units.

- 5 According to the aforementioned embodiment, packet forwarding units by which packets sent from the NIFs are forwarded to the CSW may be provided separately from packet forwarding units by which packets sent from the CSW are forwarded to the NIFs.
- 10 Hence, processing performed by the respective packet forwarding units can be limited to thereby simplify the configuration of the packet forwarding units.

- According to the aforementioned embodiment, each of the packet forwarding units may be provided to
- 15 serve as a unit for forwarding packets from the NIFs to the CSW and for forwarding packets from the CSW to the NIFs. Hence, the idle time of the respective packet forwarding units can be shortened to thereby improve the processing performance, per packet forwarding unit,
- 20 of the router.

- According to the aforementioned embodiment, a round-robin method in which packets are distributed sequentially to the packet forwarding units in entering order may be used as a packet distributing method in
- 25 the forward packet distribution units. Hence, the packet distributing method is simple and can be achieved easily. According to the aforementioned embodiment, an out-of-order method in which packets are

00767707 012401

distributed to empty ones of the packet forwarding units may be used as a packet distributing method. Hence, even in the case where the packet processing time of the packet forwarding unit varies between

5 packets, the idle time of the packet forwarding units can be shortened so that the packet forwarding performance per packet forwarding unit can be improved in view of the total performance of the router.

According to the aforementioned embodiment, a hash

10 method in which packet forwarding units used for distributing packets are univocally determined on the basis of the header conditions of the packets may be used as a packet distributing method. Hence, a series of packets to be continuously sent from the packet

15 sender address to the forwarding address are all distributed to one packet forwarding unit. Hence, overtaking of the packets can be avoided without rearranging the order of the packets before sending of the packets. Further, a pseudo-random number sequence

20 generated on the basis of the packet header conditions may be used as a hash function in the hash method so that packets are distributed to the packet forwarding units as uniformly as possible. Hence, the time deviation for the number of processed packets can be

25 suppressed among the packet forwarding units.

According to the aforementioned embodiment, the packet rearrangement units perform rearrangement of packets outputted from the packet forwarding units to

09767707.012401
T04270.20229260

make the packets rearranged in the order of packets entering the RP. Hence, overtaking of the packets can be avoided. When an out-of-order method is used as a packet distributing method, packets may be numbered sequentially when the packets are distributed. When the packets are taken around with the sequence numbers in the packet forwarding units, the packet rearrangement units can detect the entering order of the packets and can rearrange the packets in the entering order of the packets.

According to the aforementioned embodiment, there may be provided a queue waiting for the start of the packet forwarding process after reception of a packet from the forward packet distribution unit into each of the packet forwarding units, and a queue waiting for sending of packets to the forward packet rearrangement unit after the completion of processing in the packet forwarding units. Hence, the queues can absorb the deviation of the packet processing time in the respective packet forwarding units and can absorb the deviation of the time required for distributing packets to the packet forwarding units when a hash method is used as a packet distributing method in the forward packet distribution unit. Hence, the idle time of the packet forwarding units can be shortened to thereby improve the processing performance, per packet forwarding unit, of the router.

According to the aforementioned embodiment,

09767707.012401

not the packet forwarding units but the packet
retrieving units may be arranged in parallel to one
another so that a process performed by the packet
retrieving units can be limited to packet header
5 retrieval, among the processes performed by the packet
forwarding units. Hence, packet header retrieval which
is so complex as to be apt to be a bottle neck among
the processes required for every the packet can be
performed by the parallel packet retrieving units at a
10 high speed while increase in the number of hardware
devices is suppressed.

According to the aforementioned embodiment,
parallel arrangement of packet forwarding units and
parallel arrangement of packet retrieving units may be
15 executed simultaneously. Hence, the degree of
parallelism of the process required for every packet in
the router can be heightened to the degree of
parallelism of packet forwarding units and the degree
of parallelism of packet retrieving units. Hence, the
20 total performance of the router can be improved more
greatly.

According to the aforementioned embodiment,
statistical information may be picked up in each of
packet forwarding units or in each of packet retrieving
25 units. Hence, a software program managing the router
as a whole can read statistical information in each
packet forwarding unit or in each packet retrieving
unit. The information may be tabulated by the software

09767707, 012401
TOTAL 2022926

program. Alternatively, a processor may be provided in each RP so that a software program operating on the processor can read statistical information in each packet forwarding unit or in each packet retrieving unit. When the information is tabulated by the software program, the total statistical information of the router can be obtained.

According to the aforementioned embodiment, the plurality of packet forwarding units or the plurality of packet retrieving units classify packets into several groups (hereinafter referred to "flows") in accordance with the header conditions of the packets. These flows are numbered, so that flow numbers are outputted from the packet forwarding units or the packet retrieving units. The data transfer rate of the flow is measured in the forward packet rearrangement units or in the retrieval packet rearrangement unit in accordance with the flow numbers. Hence, the total data transfer rate of the flow processed by all the packet forwarding units or by all the packet retrieving units can be measured. Hence, packets can be aborted in accordance with the data transfer rate or packets can be aborted easily at the time of traffic congestion. Or the sending rate of packets can be limited when the packets are sent out.

00767707 012401
TOP SECRET